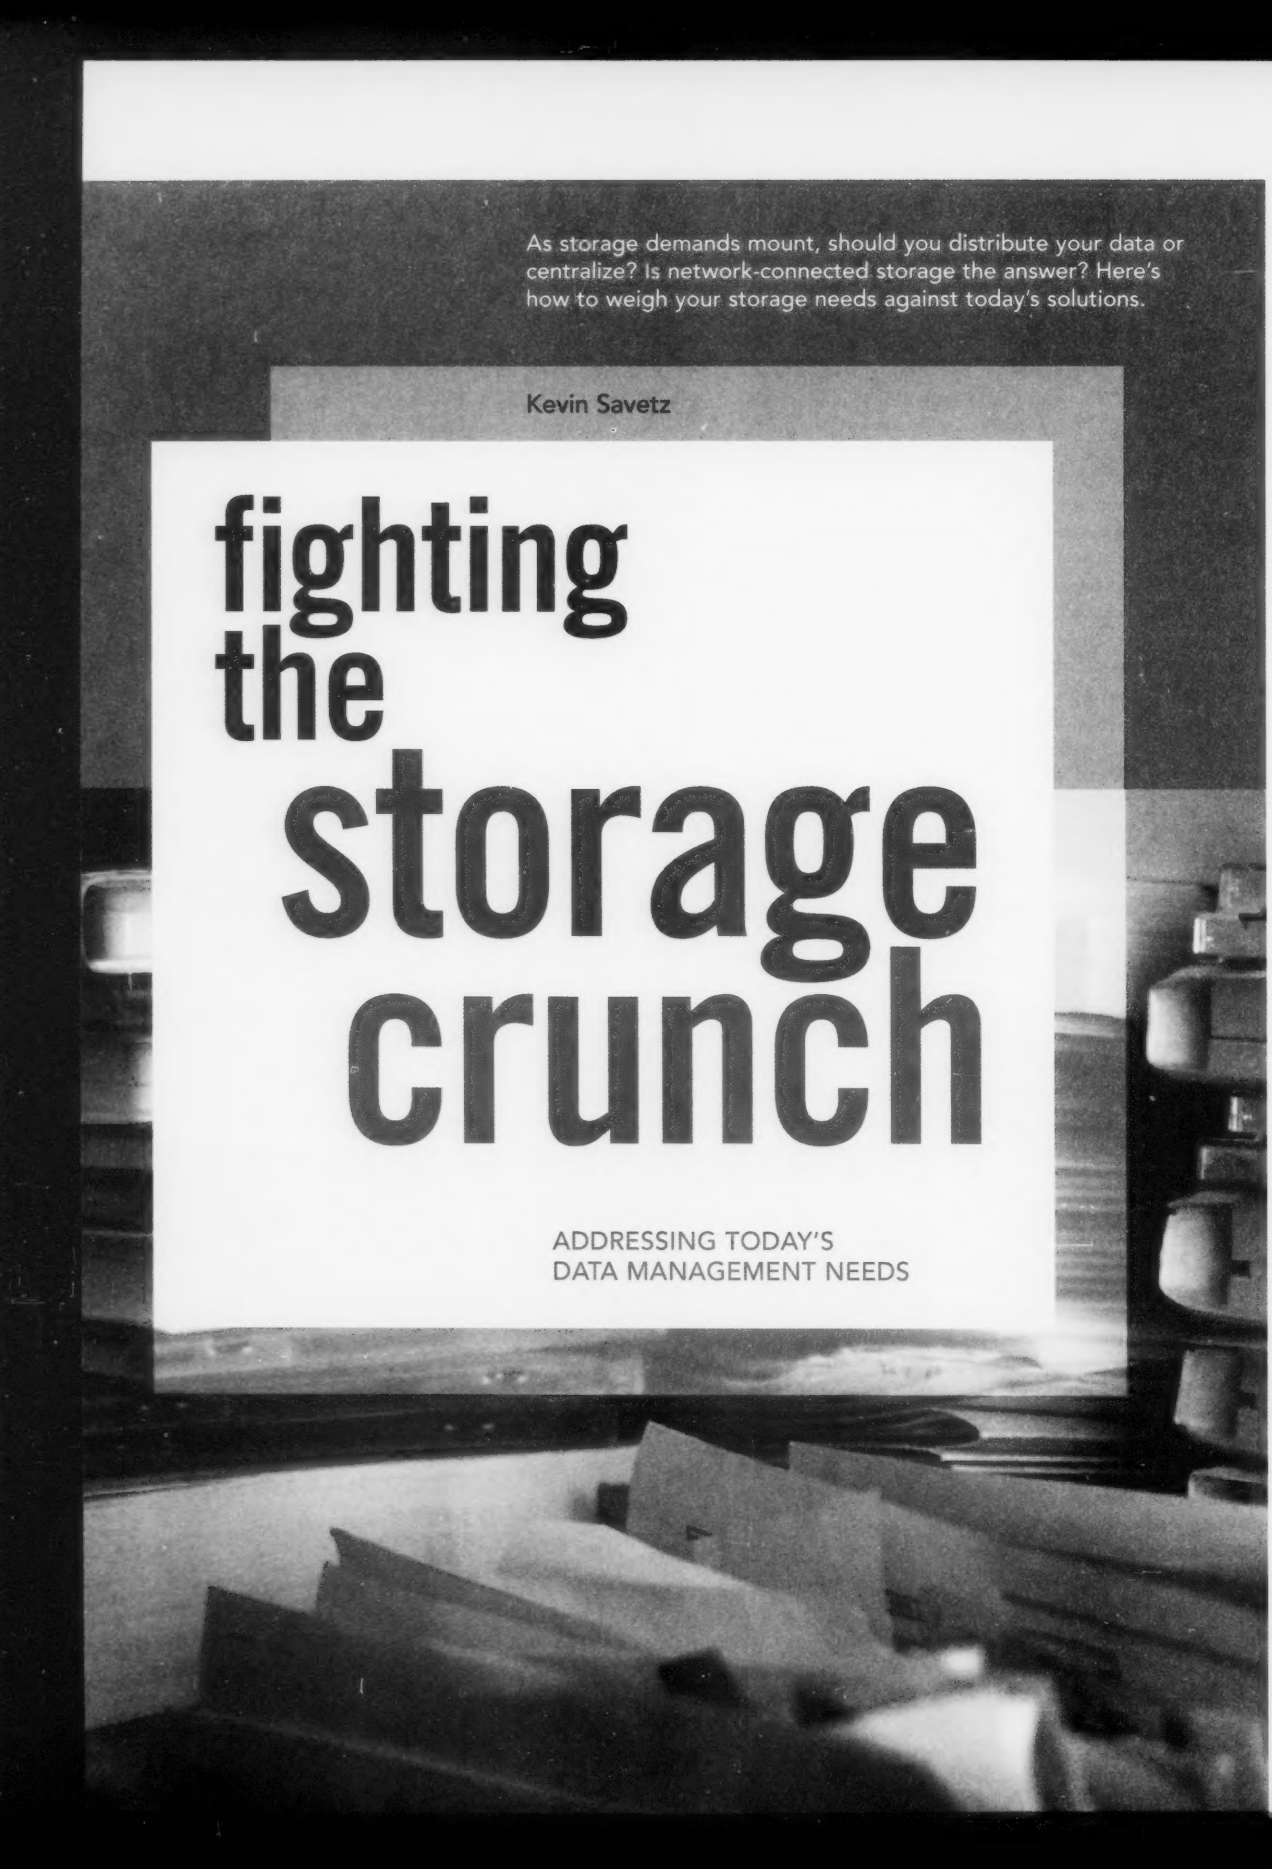

As storage demands mount, should you distribute your data or centralize? Is network-connected storage the answer? Here's how to weigh your storage needs against today's solutions.

Kevin Savetz

fighting the storage crunch

ADDRESSING TODAY'S DATA MANAGEMENT NEEDS

Photograph by Tom Falconer.

In 1958, C. Northcote Parkinson observed that "work expands so as to fill the time available for its completion." If Parkinson had worked in IT, he might have penned another law: Data expands to fill all available storage space, and then some. Demand for data storage is exploding. It's a trend that started in the 1950s and hasn't stopped; if anything, it's accelerating.

The headaches and costs associated with maintaining, managing, and protecting data increase with the capacity of storage. Backup techniques that were useful two years ago don't do the job today. Keeping garbage and obsolete data under control is harder as datasets grow. And with more data, it becomes more difficult to maintain access controls.

Performance is another perennial issue. Queries, searches, and user load all factor into the decision for the correct storage solution. Online searches are just part of the concern: Companies often need to generate complex summaries of data or perform data mining operations behind the scenes, while still making the data available to the outside world.

Choosing the right storage systems for your applications is a matter of weighing your specific needs. Speed of retrieval, the initial size of the dataset, and the anticipated growth of the dataset over time will all affect your decision. Once you've evaluated your priorities, you have a number of modern storage solutions to pick from.

a storage web

In enterprise computing, the trend is to centralize storage, allowing a small group of dedicated people to manage the preservation of data. In this model, the storage area network (SAN) has become the vehicle of choice.

A SAN is a collection of storage devices that looks like a single, very large storage device to the applications using it. The data may be stored in different locations, but the applications don't know it (and don't care). SAN vendors offer a wide variety of configurations that can integrate such devices as disk drives, tape drives, controllers, and LAN interfaces, with access controlled by a "switch" or "director."

One advantage of the SAN approach is speed. SANs can intelligently distribute requests for data, even drawing from different copies of the same data stored in different locations. They can also cache popular queries, and in some cases even perform look-ahead functions, essentially predicting what data is needed and having it ready before it's requested. Requests can be prioritized, too—so online functions can take precedence over backups, for instance.

The SAN method has its drawbacks as well. For one thing, it's difficult to maintain a redundant SAN at another location, as you might with a Web site that balances server load across multiple locations. Keeping two SANs synchronized over the Net can be difficult and bandwidth-intensive. Another disadvantage of a SAN is that you need increased network bandwidth to access the data from every part of the enterprise. For single-campus enterprises, this isn't a big deal, but if you have offices in different cities, you'll need to size your WAN links carefully.

net connected

Not all enterprises can work exclusively with the centralized storage model. In these cases, network-attached storage (NAS) can be the solution. NAS describes storage devices that connect to a network (rather than to a single server). A NAS device is actually a specialized server: a lightweight host dedicated to communicating between the storage media (usually a large hard drive or array) and other machines on the network. The initial cost for NAS is lower than for SAN, but the incremental cost may be higher.

NAS makes sense when you have a number of small sites, with each site maintained by a separate group. Security can be made much simpler by using separate drives tied together via carefully crafted URLs. It also means that the loss of one drive doesn't shut down the whole site, just the portion served by that drive.

Because the typical NAS unit has only a single-network interface, NAS can handle fewer parallel requests than SANs can. In addition, NAS controllers are usually too simple to prioritize requests intelligently, so a background task such as data mining or backup can affect online operation significantly.

For large-scale traditional transaction processing such as databases, SAN is usually the better choice because it offers consistent speed. NAS is the better option if there is much distance between the storage device and servers, or if it's necessary to share files—for instance, in software development and testing. The amount of data to be stored is typically another difference between SANs and NAS. While NAS devices might offer from 30GB to 300GB of storage—a respectable amount—SANs are often measured in terabytes.

direct connect

According to Ken Steinhardt, director of technology analysis at EMC, a leader in storage systems, NAS and SAN are both more efficient and less expensive than direct-attached storage—that is, a storage device that connects directly to a server using FireWire, USB, SCSI, IDE, or a proprietary interface. "In terms of people cost, management, operation, configuration, and backup—these can be done unbelievably more efficiently if I have the ability to consolidate with SAN or NAS," he says.

But in some cases, it makes more sense to connect a storage system to a single server. Direct-attached storage is a time-

If your business stores credit card numbers, maintaining that data independently will ensure that it's available only to those people who absolutely need it.

tested solution that can suit the storage needs of a single host well. An enterprise-level mail server might fall in this category. Using direct-connected storage for email would eliminate the network overhead of copying it to NAS or a SAN.

With manufacturers now producing hard drives that hold 160GB to 180GB, the capacity of direct-connect drives is nothing to sneeze at. A server with a RAID array configured using five of these drives has access to more than 600GB of storage. When you need more, you can add an expansion rack with another RAID set and you'll have more than a terabyte of storage that doesn't require much money or rack space.

Direct-connect storage can be beneficial even for geographically distributed multiserver Web sites. By storing static material locally, you can eliminate network traffic for those portions of the site that change very infrequently. If there's a chance that some of the locally stored pages could change, MD5 checksums are a reliable, low-bandwidth way to periodically compare local content to that hosted on the network.

Security concerns are another reason to rely on local storage rather than a network-connected solution. If your business stores credit card numbers, maintaining that data independently will ensure that it's available only to those people who absolutely need it. An independent server can also have an independent backup system, which is good, because it means the backups of confidential data can be properly controlled.

when speed is everything

Storage isn't necessarily limited to disk drives. If you need something faster, a distributed RAM disk can do the job. RAM disks have none of the seek time or rotational latency associated with traditional hard disk arrays. Let's say you have a thousand computers, each with 1.5GB of installed RAM. If you configure 1GB of that RAM on each machine to use as a distributed RAM disk, the result is a very fast one-terabyte data store.

Very few applications will require this level of storage performance, which is blindingly fast but difficult to implement properly. Such a cluster puts an extremely heavy load on networking resources, which necessitates careful design. Another downside: Running all those PCs is a huge power drain.

Because adding new RAM isn't as easy as connecting new disk drives, increasing the storage capacity of an in-memory array may require you to reorganize your data. At the terabyte level, this can be a slow and painful process, but one that can be scheduled and, perhaps, performed offline.

RAM-based storage is most effective for those portions of sites where data is changed much less often than it's read or searched. In these instances, a RAM-based cluster can serve as a front end that stores an index to data held in a larger, disk-based storage system. This way, searches will be lightning-fast, even though access to actual documents is still taking place at disk-based speeds.

mix and match

One important point to remember is that you don't have to pick just one storage solution. Indeed, a prudent system architect would use a combination of methods to implement a Web site server farm. Even the choice between NAS and SAN is not necessarily an either-or decision. IBM and HP are trying to marry SAN and NAS technology together, to reduce the time and cost of system setup without sacrificing the centralized management of storage that a pure SAN provides. If you've had problems forecasting storage needs in the past, this is well worth investigating.

Remember: With a SAN, you aren't just buying storage. You're also buying management of the storage, so choose your vendor as carefully as you would an architect.

"We see the convergence of SAN and NAS being one of the fundamental trends in the market right now," says EMC's Steinhardt. EMC offers Celerra HighRoad, software that fulfills requests for data via either a SAN or NAS, depending on which delivery method is most efficient for that request.

The key is to determine which data is stored where, and by what method the data is preserved, managed, and protected. There should be an electronic, physical, and procedural barrier between data used throughout the enterprise and data used by a small subset of the enterprise, such as credit card numbers. Determine the most cost-effective way to back up and distribute data. Finally, determine the performance you need and match it to the available technology.

landlord or architect?

You can liken data storage problems to the search for physical storage space for your business; many of the decisions and trade-offs are similar. Leasing an existing warehouse is a good analogy for network-attached storage. You sign the lease, move your stuff in, and that's that. When you outgrow it, you either lease more space, or get a larger warehouse and move out of the old one. Most importantly, you can get it today—you don't have to wait a month or more while you build your own space.

In much the same way, you can buy network-attached storage, plug it in, and you're ready to use it. Iomega offers NAS in capacities ranging from 120GB to 480GB per unit; Xtore offers larger configurations, to 1.6TB per unit; and at the top end, Land-5 offers a 20TB system that fits in a short rack. The price range on these devices varies widely, but there's something within the reach of most companies or departments.

Building a storage area network is more like buying land, erecting a building to suit your needs, then building onto it when your need for space grows. You don't just buy a SAN, you have to engineer it with your vendor, just as you engineer your building with an architect. There is lead time for a SAN, and the initial cost will be higher than for other solutions. Still, if done

properly, the integration will result in a lower total cost of ownership over time for company- or site-wide storage.

Remember: With a SAN, you aren't just buying storage. You're also buying management of the storage, so choose your vendor as carefully as you would an architect. Big vendors like IBM, HP, and StorageTek all offer SAN products, but don't overlook smaller players that may provide just the management facilities you need at lower cost. You can find more information and a partial list of vendors at the Storage Network Industry Association Web site (www.snia.org).

planning a storage strategy

For applications with moderate storage and access requirements, it's safe to stick with directly connected storage. This is especially true for secure systems, such as payment servers for e-commerce sites, and personnel information in general business where very sensitive information (such as employee medical records) is stored.

Decent rackmount servers with up to a terabyte of storage are readily available, and you can usually add an expansion satellite drive bay, should the need for even more space arise. For applications that require more storage than a single server can provide, however—such as data archiving and data mining—you'll need a larger solution.

So what is the right storage solution for your business? That depends on how much data you have and where it needs to be. In fact, your situation will probably change over time and will need to be reviewed on a regular basis.

One important point to remember is that the fundamental design of storage systems isn't likely to change over the next five years. The equipment on the market today will be what's available in the foreseeable future, with some incremental improvement. Don't put off implementing your site in hopes that something better will come along soon. Instead, focus on finding the solution that will deliver the best return for your storage buck. With costs per megabyte already at an all-time low, the only thing you can't afford to do is continue on with inadequate storage for your needs. ◘

Kevin Savetz has been a freelance technology writer for a decade. Contact him through his Web site at www.savetz.com.

desktop space

Some IT managers prefer to centralize data storage, while others are content to distribute it across geographically distant servers. But in the research facilities of a number of universities and private companies, still other developers are working on a brand new technique—one that turns the traditional approach to mass storage on its head.

A Storage Area Network (SAN) uses technologies like Fiber Channel and iSCSI to interconnect high-volume servers, appliances, and arrays to form what appears, at the application level, to be a single mass-storage device. But with the average capacities of consumer hard disks increasing at a steady rate, these large-scale repositories seldom represent the majority of an enterprise's total available storage.

Today, even the lowliest desktop workstation might have access to 40GB of storage, or more. Multiply that by the number of employees in the average workgroup, or even across the entire enterprise, and you have a formidable amount of storage space—with much of it often going unused.

Here's where peer-to-peer (P2P) technology comes in. Applications like Napster and Gnutella proved it was possible to group large numbers of anonymous workstations into a network and then retrieve information from that network, irrespective of where the information was physically housed. Now, researchers are looking for ways to transform this concept into a full-fledged persistent storage medium, one where small, low-cost computers can take the place of the expensive storage arrays that make up a traditional SAN.

The first such project was Freenet (freenetproject.org), the goal of which was to create a permanent, universally-accessible information storage network that would be impervious to filtering efforts by governments or other organizations. Since Freenet was launched in 1999, a number of other efforts have begun. One is OceanStore (oceanstore.cs.berkeley.edu), an ongoing research effort sponsored by the University of California at Berkeley. Microsoft Research has developed an experimental distributed storage system called PAST (www.research.microsoft.com/~antr/PAST/), while IBM markets distributed filesystem products upon which other P2P systems have been built.

Each of these systems has similar advantages. Additional storage can be added to P2P networks at low cost and with minimal, if any, reconfiguration of the network itself. P2P storage networks are also extremely fault tolerant—the failure of a single node doesn't bring down the whole network. But each of these systems faces common challenges, as well.

For one thing, as with any distributed computing project, P2P storage networks rely heavily on network bandwidth. While that may be fine for LANs, asynchronous broadband connections probably lack the throughput to create truly effective storage nodes. Persistence of nodes is also a problem; workstations and other general-use machines aren't generally expected to provide 24/7 access to data storage.

What's more, by their very nature, P2P networks offer little facility for centralized management. That goes for not just the nodes, but for the data itself—security is always an issue when you let sensitive information out of the data center and on to desktop workstations.

Developers are well aware of these concerns, however. As accompanying technologies mature, like encryption systems and redundant storage algorithms, P2P storage networks are likely to move away from the fringes of computing to become an attractive option for enterprise-wide mass storage.

—Neil McAllister